

GRADIENT ESTIMATION FOR RATIOS

AD-A248 453



by

Peter W. Glynn, Pierre L'Ecuyer and Michel Adès

TECHNICAL REPORT No. 70

AUGUST 1991

Prepared under the Auspices

of

U.S. Army Research Contract

DAAL-03-91-G-0101

Approved for public release: distribution unlimited.

Reproduction in whole or in part is permitted for any
purpose of the United States government.

DEPARTMENT OF OPERATIONS RESEARCH
STANFORD UNIVERSITY
STANFORD, CALIFORNIA

98 4 10 039

92-09252



DTIC
ELECTE
APR 13 1992
S D D



GRADIENT ESTIMATION FOR RATIOS

Peter W. Glynn*
Operations Research Department
Stanford University
Stanford, CA 94305-4022, USA

Pierre L'Ecuyer
Département d'IRO
Université de Montréal
C.P. 6128, Succ. A
Montréal, H3C 3J7, Canada

Michel Adès
Département de Math. et Info.
Université du Québec à Montréal
C.P. 8888, Succ. A
Montréal, H3C 3P8, Canada

ABSTRACT

The ratio estimation problem arises in many different applications settings. This paper is concerned with the interplay between gradient estimation and ratio estimation. Given unbiased estimators for the numerator and the denominator of the ratio, as well as their gradients, joint central-limit theorems for the ratio and its gradient are derived. The resulting confidence regions are of potential interest when optimizing such ratios numerically, or for sensitivity analysis with respect to parameters whose exact value is unknown. The paper also briefly discusses low-bias estimation for the gradient of a ratio.

1. INTRODUCTION

Let (A, B) be a pair of jointly distributed real-valued random variables. The estimation of the ratio $\alpha = E[A]/E[B]$ is known, in the simulation literature, as the *ratio estimation problem*. Such ratio estimation problems arise in many different applications settings. For example, it is well known that the steady-state mean of a positive recurrent regenerative stochastic process can be expressed as such a ratio of expectations; see, for example, Section 3.3.2 of Bratley, Fox, and Schrage (1987) or Chapter 2 of Wolff (1989). In Section 2 of this paper, we will discuss the ratio estimation problem in greater detail and offer additional examples. It will turn out that the infinite-horizon discounted cost of a non-delayed regenerative process can also be expressed in terms of an appropriately chosen ratio estimation problem. This fact was first pointed out by Fox and Glynn (1989).

Recently, the simulation community has devoted a great deal of attention to the use of simulation as an optimization tool. An important component of this research effort has been the development of estimation methodology for computing the gradient of a real-valued performance measure with respect to a (finite-dimensional) decision parameter vector. Such gradients play an important role in many iterative algorithms for performing both constrained and unconstrained mathematical optimization. This paper is intended as a study of the question of how to use this gradient estimation methodology

in the setting of the ratio estimation problem.

The paper is organized as follows. In Section 2, a number of different applications in which ratio estimation problems arise are discussed, and the mathematical framework for the remainder of the paper is described. Section 3 is devoted to deriving a confidence interval methodology for estimating the partial derivative of a ratio. In addition, a joint central-limit theorem for the simultaneous estimation of the entire gradient is obtained. In Section 4, low-bias estimation issues are discussed. Finally, Section 5 discusses some experimental results related to gradient estimators for ratios, and Section 6 concludes the paper with a brief summary. The proof of our main theorem (Theorem 1) is given in the Appendix. The other proofs are not given here. A (future) more elaborate version of the paper will contain all the proofs, derive a joint central-limit theorem that can be used to simultaneously estimate the gradient and the Hessian of mixed second-partial derivatives of a ratio, and provide further numerical illustrations.

2. EXAMPLES OF RATIO ESTIMATION PROBLEMS

As discussed in the introduction, the ratio estimation problem is concerned with the estimation of the ratio

$$\alpha = \frac{E[A]}{E[B]},$$

where (A, B) is a pair of jointly distributed real-valued random variables. We now proceed to offer several examples of this estimation problem.

EXAMPLE 1. Let $X = \{X(t), t \geq 0\}$ be a real-valued (possibly) delayed regenerative process with regenerative times $0 \leq T(0) < T(1) < \dots$. For $i \geq 1$, let

$$\begin{aligned}\bar{A}_i &= \int_{T(i-1)}^{T(i)} |X(s)| ds \\ A_i &= \int_{T(i-1)}^{T(i)} X(s) ds \\ B_i &= T(i) - T(i-1).\end{aligned}$$

*Research supported by the U.S. Army Research Office under Contract DAAL-03-91-G-0101.

If $E[\bar{A}_1 + B_1] < \infty$, then it can be shown (see, for example, Asmussen 1987, or Wolff 1989) that

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t X(s) ds = \alpha = E[A_1]/E[B_1].$$

Hence, as discussed in the introduction, the steady-state mean of such a process can be expressed as the ratio of the two expectations $E[A_1]$ and $E[B_1]$.

EXAMPLE 2. Let $X = \{X(t), t \geq 0\}$ be a non-delayed regenerative process, taking values in a state space S , with regenerative times $0 = T(0) < T(1) < \dots$. Let f and g be two real-valued non-negative (measurable) functions defined on S , and set

$$V(t) = \int_0^t g(X(s)) ds$$

$$\alpha = E \left[\int_0^\infty \exp[-V(t)] f(X(t)) dt \right].$$

Then, α is the infinite-horizon expected discounted cost, the process $g(X(t))$ corresponds to the (state-dependent) discount rate at time t , and $f(X(t))$ is the (undiscounted) rate at which cost is incurred at time t . A common choice for g is the one in which $g(\cdot)$ is constant and equal to $\rho > 0$, in which case

$$\alpha = E \left[\int_0^\infty \exp[-\rho t] f(X(t)) dt \right]$$

is the infinite-horizon ρ -discounted cost. Let

$$A_1 = \int_{T(0)}^{T(1)} \exp \left[- \int_0^t g(X(s)) ds \right] f(X(t)) dt$$

$$C_1 = \exp [-V(T(1))]$$

$$B_1 = 1 - C_1.$$

Because of the regenerative structure of X , it is evident that α satisfies the equation $\alpha = E[A_1] + E[C_1]\alpha$. Thus, if $E[C_1] < 1$, it follows that α is finite and can be expressed as

$$\alpha = \frac{E[A_1]}{E[B_1]}.$$

Hence, the infinite-horizon discounted cost for a regenerative process can be expressed in terms of a ratio estimation problem; see Fox and Glynn (1989) for further details.

EXAMPLE 3. Let X be a regenerative process as in Example 2, and assume that X has right-continuous paths with left limits. Let F be a non-empty subset of the state space S , and let $\tau(F) = \inf\{t \geq 0 \mid X(t) \in F\}$ be the first hitting time of the subset F . Then,

$$\alpha = E[\tau(F)]$$

is the mean hitting time of F . Such expectations are of interest, for example, in the reliability setting, in which

case $\tau(F)$ would typically correspond to the system failure time, and $T(1)$ to a time at which the system is brought back to an "as good as new" state. Let

$$A_1 = \min[\tau(F), T(1)]$$

$$B_1 = I[\tau(F) < T(1)],$$

where I denotes the indicator function. If $P[\tau(F) < \infty] > 0$ (note that this is equivalent to requiring that $P[\tau(F) < T(1)] > 0$), it is easily shown that

$$\alpha = \frac{E[A_1]}{E[B_1]}.$$

See Goyal et al. (1991) for additional details. Thus, the mean hitting time of a regenerative process can be formulated in terms of the ratio estimation problem.

EXAMPLE 4. Let X be a real-valued random variable and let C be an event with $P(C) > 0$. Suppose that we wish to estimate

$$\alpha = E[X \mid C],$$

namely the conditional expectation of X , given that the event C has occurred. If $E[|X|] < \infty$, then we can express α in terms of the ratio $\alpha = E[A_1]/E[B_1]$, where

$$A_1 = X I(C)$$

$$B_1 = I(C).$$

Hence, conditional expectations are expressible in terms of the ratio estimation problem.

Thus, the ratio estimation problem arises in a variety of different applications contexts. We shall now introduce a decision parameter vector θ into the discussion. For each $\theta \in \mathbb{R}^d$, let P_θ be the probability measure associated with the parameter value θ , and let E_θ be its corresponding expectation operator. In addition, we shall permit the random variables $A(\theta)$ and $B(\theta)$ to depend explicitly on $\theta \in \mathbb{R}^d$. Then, for each $\theta \in \mathbb{R}^d$, the ratio of expectations can be expressed in the form

$$\alpha(\theta) = \frac{u(\theta)}{\ell(\theta)},$$

where $u(\theta) = E_\theta[A(\theta)]$ and $\ell(\theta) = E_\theta[B(\theta)]$. Given our above examples, computing the gradient of such a ratio $\alpha(\theta)$ is useful for sensitivity analysis or optimization of any of the following: steady-state costs or rewards in regenerative processes; infinite-horizon discounted costs; mean time to failure in reliability systems; conditional expectations and probabilities.

3. CONFIDENCE INTERVALS FOR GRADIENT ESTIMATORS OF RATIOS

Let $\theta_0 \in \mathbb{R}^d$ be fixed. In order for the gradient estimation problem to make sense, we shall require that both

$u(\cdot)$ and $\ell(\cdot)$ have gradients at $\theta = \theta_0$. We shall further assume that there exists unbiased estimators for not only $u(\theta_0)$ and $\ell(\theta_0)$, but also their gradients $\nabla u(\theta_0)$ and $\nabla \ell(\theta_0)$. Focussing now on the i -th component of the gradient, we shall specifically assume that there exist jointly distributed random variables (A, B, C, D) such that

$$\begin{aligned} E[A] &= u(\theta_0) \\ E[B] &= \ell(\theta_0) \\ E[C] &= \partial_{\theta_1} u(\theta_0) \stackrel{\text{def}}{=} \left. \frac{\partial}{\partial \theta_1} u(\theta) \right|_{\theta = \theta_0} \\ E[D] &= \partial_{\theta_1} \ell(\theta_0) \stackrel{\text{def}}{=} \left. \frac{\partial}{\partial \theta_1} \ell(\theta) \right|_{\theta = \theta_0} \end{aligned}$$

where ∂_i denotes the partial derivative with respect to θ_i , and θ_i is the i -th component of θ .

There is now a great deal of literature on various ways of constructing unbiased estimators for $\partial_t u(\theta_0)$ and $\partial_t \ell(\theta_0)$. The two principal approaches that have been explored are likelihood ratio gradient estimation (see Glynn 1990 for a survey) and infinitesimal perturbation analysis (see Glasserman 1991). For links between the two methods and for a general survey, see L'Ecuyer (1990, 1991).

We shall now assume that it is possible for the simulator to generate a sequence $\{(A_j, B_j, C_j, D_j), j \geq 1\}$ of i.i.d. replicates of the random vector (A, B, C, D) . In each of the problem settings described in Section 2, this is typically straightforward.

To estimate

$$\begin{aligned}\partial_{\alpha}\ell(\theta_0) &= \frac{\ell(\theta_0)\partial_{\alpha}u(\theta_0) - u(\theta_0)\partial_{\alpha}\ell(\theta_0)}{\ell^2(\theta_0)} \\ &= \frac{\partial_{\alpha}u(\theta_0) - \alpha(\theta_0)\partial_{\alpha}\ell(\theta_0)}{\ell(\theta_0)},\end{aligned}$$

the natural estimator to use is

$$\delta_i(n) = \frac{C_n - \alpha_n D_n}{\beta_n},$$

where

$$\begin{aligned} A_n &= \frac{1}{n} \sum_{j=1}^n A_j \\ B_n &= \frac{1}{n} \sum_{j=1}^n B_j \\ C_n &= \frac{1}{n} \sum_{j=1}^n C_j \\ D_n &= \frac{1}{n} \sum_{j=1}^n D_j \end{aligned}$$

and

$$\alpha_n = \lambda_n / B_n.$$

Our first proposition states that under reasonable conditions, $\delta_1(n)$ is a consistent estimator for $\partial_1 \alpha(\theta_0)$. The proof is straightforward and therefore omitted.

PROPOSITION 1. *Suppose that $E[|A_1| + |B_1| + |C_1| + |D_1|] < \infty$ and that $E[B_1] \neq 0$. Then,*

$$\lim_{n \rightarrow \infty} \delta_i(n) \stackrel{a.s.}{=} \partial_i \alpha(\theta_0). \quad \blacksquare$$

To develop a confidence interval methodology for $\partial_1(n)$, we need a central-limit theorem (CLT) for the estimator. Let

$$\begin{aligned} Z_j &= A_j - \alpha(\theta_0)B_j, \\ W_j &= C_j - \alpha(\theta_0)D_j - \partial_i \alpha(\theta_0)B_j, \end{aligned}$$

and note that under the assumptions of Proposition 1, $E[Z_j] = E[W_j] = 0$. This observation is an important element in the proof of the following theorem.

THEOREM 1. Assume that $E[Z_1^2 + W_1^2] < \infty$. If, in addition, the conditions of Proposition 1 are in force, then

$$\sqrt{n}[\delta_i(n) - \partial_i \alpha(\theta_0)] \Rightarrow \sigma N(0, 1)$$

as $n \rightarrow \infty$, where

$$\sigma^2 = \frac{E[W_1 - (E[D_1]/E[B_1])Z_1]^2}{(E[B_1])^2}.$$

Theorem 1 has been previously established, using different methods, by Reiman and Weiss (1989) in the context of likelihood ratio gradient estimation for regenerative steady-state simulation. Their expression for the variance constant σ^2 is formally different, but algebraically identical.

The final step needed to develop a confidence interval methodology for $\delta_i(n)$ is the construction of an appropriate estimator for σ^2 . Let

$$v(n) = \frac{\frac{1}{n} \sum_{j=1}^n [W_j - (D_n/B_n)Z_j]^2}{(B_n)^2}$$


where

$$\begin{aligned} Z_j &= A_j - \alpha_n B_j, \\ W_j &= C_j - \alpha_n D_j - \delta_j(n) B_j. \end{aligned}$$

The next proposition gives conditions under which $v(n)$ is strongly consistent for σ^2 . The proof is straightforward and therefore omitted.

PROPOSITION 2. Suppose that $E[A_1^2 + B_1^2 + C_1^2 + D_1^2] < \infty$. If $E[B_1] \neq 0$, then

$$\lim_{n \rightarrow \infty} v(n) \stackrel{\text{a.s.}}{=} \sigma^2. \blacksquare$$



A-1

We note that if $v(n)$ is computed via a two-pass approach in which α_n and $\delta_i(n)$ are computed in the first pass through the data $\{(A_j, B_j, C_j, D_j), 1 \leq j \leq n\}$ and the sum of squares computed in the second pass, then it is essentially guaranteed that $v(n)$ will be computed as a non-negative quantity on any finite-precision computer. More importantly, this means of computing $v(n)$ is likely to be more stable numerically than that associated with the computation described in Reiman and Weiss (1989).

We are now ready to describe a general confidence interval methodology for estimating partial derivatives of ratios. Suppose that we wish to compute a $100(1 - \delta)\%$ confidence interval for $\partial_i \alpha(\theta_0)$. We use the following procedure:

Algorithm CI.

1. Generate $\{(A_j, B_j, C_j, D_j), j \geq 1\}$.
2. Compute α_n and $\delta_i(n)$.
3. Compute $v(n)$ (using the two-pass approach described above).
4. Find $z(\delta)$ such that $P[N(0, 1) \leq z(\delta)] = 1 - \delta/2$.
5. Compute

$$\begin{aligned} L_n &= \delta_i(n) - z(\delta) \sqrt{v(n)/n} \\ U_n &= \delta_i(n) + z(\delta) \sqrt{v(n)/n}. \blacksquare \end{aligned}$$

Then, $[L_n, U_n]$ is an (approximate) $100(1 - \delta)\%$ confidence interval for $\partial_i \alpha(\theta_0)$. In particular, if the conditions of Proposition 2 are in force and $\sigma^2 > 0$, then

$$\lim_{n \rightarrow \infty} P[\partial_i \alpha(\theta_0) \in [L_n, U_n]] = 1 - \delta.$$

We conclude this section with a brief discussion of the problem of generating a confidence region for the vector $(\alpha(\theta_0), \partial_1 \alpha(\theta_0), \dots, \partial_d \alpha(\theta_0))$. A joint confidence region could be of potential interest in a number of optimization settings, since virtually all iterative (deterministic) optimization algorithms choose their search direction, at each iteration, by considering the full gradient.

Let $C(i)$ and $D(i)$ be unbiased estimators for $\partial_i u(\theta_0)$ and $\partial_i \ell(\theta_0)$, so that

$$\begin{aligned} E[C(i)] &= \partial_i u(\theta_0) \\ E[D(i)] &= \partial_i \ell(\theta_0). \end{aligned}$$

If $\{(A_j, B_j, C_j(1), D_j(1), \dots, C_j(d), D_j(d)), 1 \leq j \leq n\}$ is a set of n i.i.d. replicates of the random vector $(A, B, C(1), D(1), \dots, C(d), D(d))$, then the estimators $\alpha_n, \delta_1(n), \dots, \delta_d(n)$ can be constructed from the sample in the obvious way, namely

$$\begin{aligned} \alpha_n &= A_n/B_n \\ \delta_i(n) &= (C_n(i) - \alpha_n D_n(i))/B_n. \end{aligned}$$

Define

$$W_j(i) = C_j(i) - \alpha(\theta_0) D_j(i) - \partial_i \alpha(\theta_0) B_j.$$

We are now ready to state a joint CLT for $(\alpha_n, \delta_1(n), \dots, \delta_d(n))$.

THEOREM 2. Assume that $E[A_1^2 + B_1^2 + C_1^2(1) + D_1^2(1) + \dots + C_1^2(d) + D_1^2(d)] < \infty$. If $E[B_1] \neq 0$, then

$$\begin{aligned} \sqrt{n}[\alpha_n - \alpha(\theta_0), \delta_1(n) - \partial_1 \alpha(\theta_0), \dots, \\ \delta_d(n) - \partial_d \alpha(\theta_0)] E[B_1] \Rightarrow N(0, C) \end{aligned}$$

as $n \rightarrow \infty$, where $C = (C_{ij}, 0 \leq i, j \leq d)$ is a covariance matrix whose elements are given by

$$\begin{aligned} C_{00} &= E[Z_1^2] \\ C_{0i} &= C_{i0} \\ &= E\left[\left(W_1(i) - \frac{E[D_1(i)]}{E[B_1]} Z_1\right) Z_1\right] \\ C_{ij} &= C_{ji} \\ &= E\left[\left(W_1(i) - \frac{E[D_1(i)]}{E[B_1]} Z_1\right) \left(W_1(j) - \frac{E[D_1(j)]}{E[B_1]} Z_1\right)\right] \end{aligned}$$

for $1 \leq i, j \leq d$. ■

The proof of this theorem mirrors that of Theorem 1 and is therefore omitted.

A procedure for producing asymptotically valid confidence regions for $(\alpha(\theta_0), \partial_1 \alpha(\theta_0), \dots, \partial_d \alpha(\theta_0))$ can now easily be derived, using the same ideas as those described earlier in this section for $\partial_i \alpha(\theta_0)$. We leave the details to the reader.

4. LOW BIAS ESTIMATION FOR THE GRADIENT OF A RATIO

Since the gradient of the ratio is a nonlinear function of the expectations $E[A]$, $E[B]$, $E[C(1)]$, $E[D(1)]$, \dots , $E[C(d)]$, $E[D(d)]$, it follows that the estimator $\delta_i(n)$ is, in general, biased for $\partial_i \alpha(\theta_0)$.

We will now proceed to (formally) derive a bias expansion for $\delta_i(n)$. The proof of Theorem 1 shows that

$$\delta_i(n) - \partial_i \alpha(\theta_0) = \frac{W_n - (D_n/B_n)Z_n}{B_n}. \quad (1)$$

We would like to approximate the expectation of that. We note that since B_n is close to $\mu \stackrel{\text{def}}{=} E[B_1]$ for large n , we can use the power series expansion for $f(x) = (1 - x)^{-1}$ to obtain

$$\frac{1}{B_n} = \frac{1}{\mu} \left[1 - \left(1 - \frac{B_n}{\mu} \right) \right]^{-1}$$

$$\begin{aligned}
&= \mu^{-1} \left[1 + \left(1 - \frac{B_n}{\mu} \right) + \left(1 - \frac{B_n}{\mu} \right)^2 + \dots \right] \\
&\approx \mu^{-1} \left[1 + \left(1 - \frac{B_n}{\mu} \right) \right] \\
&= \frac{2\mu - B_n}{\mu^2}.
\end{aligned}$$

Using this approximation in (1), we find that

$$\begin{aligned}
&\delta_i(n) - \partial_i \alpha(\theta_0) \\
&\approx \frac{2\mu - B_n}{\mu^2} \left(W_n - \frac{Z_n D_n}{B_n} \right) \\
&= \frac{2W_n}{\mu} + \frac{Z_n D_n - B_n W_n}{\mu^2} - \frac{2Z_n D_n}{\mu B_n} \\
&\approx \frac{2W_n}{\mu} + \frac{Z_n D_n - B_n W_n}{\mu^2} - \frac{2Z_n D_n(2\mu - B_n)}{\mu^3} \\
&= \frac{2W_n}{\mu} - \frac{Z_n D_n + B_n W_n}{\mu^2} + \frac{2Z_n D_n B_n}{\mu^3}, \quad (2)
\end{aligned}$$

where $\hat{B}_n = B_n - \mu$. Recall that $E[W_j] = E[Z_j] = 0$. Observe that for $i \neq j$, $E[B_i W_j] = E[B_i]E[W_j] = 0$, since B_i and W_j are independent. Therefore,

$$E[B_n W_n] = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n E[B_i W_j] = \frac{E[B_1 W_1]}{n}.$$

Similarly, $E[Z_n D_n] = E[Z_1 D_1]/n$. Also, $E[Z_i D_j (B_k - \mu)] = 0$ whenever $i \neq k$. Therefore,

$$\begin{aligned}
E[Z_n D_n \hat{B}_n] &= \frac{E[Z_1 D_1 (B_1 - \mu)]}{n^2} \\
&\quad + \frac{(n-1)E[Z_1 (B_1 - \mu)]E[D_1]}{n^2} \\
&= \frac{E[Z_1 B_1]E[D_1]}{n} + o(1/n).
\end{aligned}$$

Now, taking the expectation in (2) yields

$$\begin{aligned}
E[\delta_i(n)] - \partial_i \alpha(\theta_0) \\
&\approx \frac{2E[Z_1 B_1]E[D_1] - \mu E[B_1 W_1 + Z_1 D_1]}{n\mu^3}.
\end{aligned}$$

This bias approximation suggests an obvious means of reducing the bias of gradient estimators for ratios. The idea is to estimate the bias term and correct for it by subtracting off the estimated bias. In this case, this approach leads to the estimator

$$\begin{aligned}
\tilde{\delta}_i(n) &= \delta_i(n) + \frac{2D_n \sum_{j=1}^n \hat{Z}_j B_j}{n^2 \hat{B}_n^2} \\
&\quad - \frac{\sum_{j=1}^n (B_j W_j + \hat{Z}_j D_j)}{n^2 \hat{B}_n^2},
\end{aligned}$$

where \hat{Z}_j and W_j are defined just before the statement of Proposition 2 in Section 3.

Under appropriate regularity hypotheses, and by applying techniques similar to those used in Glynn and Heidelberger (1991), one can rigorously prove that $\tilde{\delta}_i(n)$ reduces the asymptotic bias, in the sense that

$$E[\tilde{\delta}_i(n)] = \partial_i \alpha(\theta_0) + o(1/n).$$

A second approach that is frequently used to correct for "nonlinearity bias" of the above type is to "jackknife" the estimator. Specifically, for $1 \leq j \leq n$, let

$$\begin{aligned}
\alpha_{n(j)} &= \frac{\sum_{k=1, k \neq j}^n A_k}{\sum_{k=1, k \neq j}^n B_k} \\
\beta_{n(j)} &= \frac{\left(\sum_{k=1, k \neq j}^n C_k \right) - \alpha_{n(j)} \sum_{k=1, k \neq j}^n D_k}{\sum_{k=1, k \neq j}^n B_k} \\
\delta_{n(j)} &= n\alpha_{n(j)} - (n-1)\beta_{n(j)}.
\end{aligned}$$

Then,

$$\delta_i^j(n) = \frac{1}{n} \sum_{j=1}^n \delta_{n(j)}$$

is the jackknife estimator for $\partial_i \alpha(\theta_0)$. Also,

$$\sqrt{n} \frac{(\delta_i^j(n) - \partial_i \alpha(\theta_0))}{s^j(n)} \Rightarrow N(0, 1),$$

where

$$(s^j(n))^2 = \frac{1}{n-1} \sum_{j=1}^n (\delta_{n(j)} - \delta_i^j(n))^2$$

is a consistent variance estimator. As in the case of the estimator $\tilde{\delta}_i(n)$, one can prove rigorously (under suitable regularity hypotheses) that the estimator $\delta_i^j(n)$ reduces asymptotic bias, in the sense that

$$E[\delta_i^j(n)] = \partial_i \alpha(\theta_0) + o(1/n).$$

It turns out that the improved bias characteristics of these estimators are costless relative to the variance, in the sense that the estimators $\tilde{\delta}_i(n)$ and $\delta_i^j(n)$ obey precisely the same CLT as does $\delta_i(n)$. Hence, the estimators exhibit the same degree of asymptotic variability.

THEOREM 3. Assume that $E[A_1^2 + B_1^2 + C_1^2 + D_1^2] < \infty$ and that $E[B_1] \neq 0$. Then,

$$\sqrt{n}(\tilde{\delta}_i(n) - \partial_i \alpha(\theta_0)) \Rightarrow \sigma N(0, 1)$$

$$\sqrt{n}(\delta_i^j(n) - \partial_i \alpha(\theta_0)) \Rightarrow \sigma N(0, 1)$$

where σ^2 is the same constant as in Theorem 1. ■

Table 1: Bias, Half-Widths, and Coverages of Confidence Intervals for $\theta = 0.2$
($\alpha(\theta) = 0.25$ and $\alpha'(\theta) = 1.5625$)

	$n = 10$			$n = 100$			$n = 1000$		
	bias	half-width	cover.	bias	half-width	cover.	bias	half-width	cover.
α_n	-.010±.002	.129±.002	.77±.01	-.001±.001	.059±.001	.90±.01	-.000±.001	.021±.001	.94±.01
α_n^J	-.001±.002	.149±.002	.80±.01	-.000±.001	.061±.001	.90±.01	-.000±.001	.021±.001	.94±.01
$\delta(n)$	-.401±.024	.914±.021	.43±.01	-.113±.013	.857±.014	.74±.01	-.004±.005	.391±.004	.88±.01
$\delta^J(n)$	-.120±.035	1.345±.038	.52±.01	-.014±.013	.921±.017	.76±.01	-.001±.005	.395±.004	.88±.01

Table 2: Bias, Half-Widths, and Coverages of Confidence Intervals for $\theta = 0.5$
($\alpha(\theta) = 1.0$ and $\alpha'(\theta) = 4.0$)

	$n = 10$			$n = 100$			$n = 1000$		
	bias	half-width	cover.	bias	half-width	cover.	bias	half-width	cover.
α_n	-.135±.009	.440±.005	.57±.01	-.019±.004	.313±.004	.82±.01	-.003±.002	.125±.001	.92±.01
α_n^J	-.036±.011	.534±.012	.65±.01	-.001±.004	.339±.004	.84±.01	-.001±.002	.126±.001	.92±.01
$\delta(n)$	-2.072±.050	1.692±.057	.26±.01	-.407±.048	2.333±.051	.57±.01	-.056±.019	1.493±.022	.81±.01
$\delta^J(n)$	-1.152±.089	2.694±.084	.37±.01	-.049±.063	2.919±.077	.63±.01	-.012±.020	1.541±.022	.82±.01

5. A NUMERICAL EXAMPLE

We will now illustrate some of the ideas developed in this paper with a numerical example. We consider the steady-state sojourn time of a customer in an M/M/1 queue with arrival rate $\lambda = 1$ and mean service time θ , where $0 < \theta < 1$. The sojourn time is the sum of the time spent by a customer waiting for service, plus that customer's service time. Let $X_0 = 0$ and, for $i \geq 1$, let X_i be the sojourn time of the i -th customer, starting from an empty system. It is well known that

$$X_i := (X_{i-1} - \nu_i)^+ + \theta\zeta_i,$$

where $\{\nu_1, \zeta_1, \nu_2, \zeta_2, \dots\}$ is a sequence of i.i.d. $\exp(1)$ random variables. For this model, the mean steady-state sojourn time $\alpha(\theta)$ can be computed in closed form:

$$\alpha(\theta) = \theta/(1 - \theta).$$

Hence,

$$\alpha'(\theta) = 1/(1 - \theta)^2$$

This system regenerates when customers arrive to an empty queue. Consequently, as discussed in Example 1 of Section 2, the steady-state mean sojourn time of a customer can be expressed in terms of a ratio estimation problem, and the methodology of this paper is therefore applicable. It is also straightforward to apply the likelihood ratio method for gradient estimation to this problem

(see L'Ecuyer, Giroux, and Glynn 1991), thereby obtaining the required unbiased estimators for the numerator and denominator of the ratio (as discussed in Section 3). It turns out that while infinitesimal perturbation analysis can be applied to obtain strongly consistent steady-state gradient estimators for this problem, it fails to give unbiased estimators of the gradient of the numerator and of the gradient of the denominator of the regenerative ratio formula. See Heidelberger et al. (1988) for further details. As a consequence, the theory of this paper is not applicable to the infinitesimal perturbation analysis steady-state gradient estimator for this problem. But on the other hand, the infinitesimal perturbation analysis derivative estimator is itself the estimator of a ratio of expectations, so that one can apply the standard theory relative to the construction of confidence intervals for ratios of expectations (Iglehart 1975, Wolff 1989).

Tables 1 and 2 report the experimental results obtained for this example. Simulation runs were carried out at two parameter values, namely $\theta = 0.2$ and $\theta = 0.5$, using $n = 10, 100$, and 1000 regenerative cycles. A total of four estimators were considered in this experiment, namely the ratio estimator α_n for $\alpha(\theta)$ and its jackknifed analog α_n^J (see Iglehart 1975), and the derivative estimator $\delta(n)$ (Section 3) for $\alpha'(\theta)$ and its corresponding jackknifed analog $\delta^J(n)$ (Section 4). Standard regenerative confidence intervals were constructed for the estimator α_n , and the confidence interval methodology of Section 3 was used to analyze $\delta(n)$. For the jackknifed versions,

confidence intervals were constructed based on the variance estimator $(s^J(n))^2$ given in Section 4. At each of the two parameter values and three choices of n (the number of regenerative cycles), a total of 10,000 95% confidence intervals was replicated for each of the four estimators. From that, we are able to report estimates for the bias, expected half-width, and coverage (the probability that the quantity being estimated lies in the confidence interval), as well as 95% confidence intervals for the bias, expected half-width, and coverages themselves.

One can observe that for small n , for all estimators, the coverage is really lower than what is to be expected. This bad behavior gets worse when θ increases (heavier traffic). Jackknifing clearly reduces the bias significantly. It also gives a better coverage for small n , but usually at the expense of a wider confidence interval. For small n , the coverage is too low anyway. For larger n , jackknifing still helps reducing the bias, but (perhaps surprisingly) does not improve the coverage significantly. Of course, this is just a particular illustration, and one must be careful about drawing any general conclusions from these numerical results.

6. CONCLUSION

Ratio estimation problems arise in many different applications settings. When estimation is to be used to analyze the sensitivity of (or to optimize) a system in which the ratio estimation problem occurs, the results of this paper become pertinent. We have derived a numerically stable confidence interval procedure for computing partial derivatives of such ratios, and have developed the appropriate joint CLT's necessary to extend this methodology to the computation of confidence regions for the full gradient of the ratio. In addition, we have discussed low-bias estimators for computing such partial derivatives. We have also described preliminary computational experience with some of the methods developed in this paper.

ACKNOWLEDGMENTS

The work of the first author was supported by the U.S. Army Research Office under Contract DAAL03-91-G-0101. The work of the second author was supported by NSERC-Canada grant no. OGP0110050 and FCAR grant no. EQ2831.

APPENDIX

Proof of Theorem 1. We note that

$$\begin{aligned} & B_n[\delta_i(n) - \partial_i \alpha(\theta_0)] \\ &= C_n - \alpha_n D_n - \partial_i \alpha(\theta_0) B_n \\ &= W_n - (\alpha_n - \alpha(\theta_0)) D_n \\ &= W_n - (D_n/B_n) Z_n \end{aligned}$$

$$\begin{aligned} &= W_n - (E[D_1]/E[B_1]) Z_n \\ &\quad - (D_n/B_n - E[D_1]/E[B_1]) Z_n. \end{aligned}$$

Clearly, $\sqrt{n} Z_n \Rightarrow (E[Z_1^2])^{1/2} N(0, 1)$ as $n \rightarrow \infty$ and $D_n/B_n \xrightarrow{a.s.} E[D_1]/E[B_1]$ as $n \rightarrow \infty$. It follows, by the converging-together principle, that

$$\sqrt{n} (D_n/B_n - E[D_1]/E[B_1]) Z_n \Rightarrow 0$$

as $n \rightarrow \infty$. The CLT for i.i.d. random variables also proves that

$$\sqrt{n} (W_n - (E[D_1]/E[B_1]) Z_n) \Rightarrow E[B_1] \sigma N(0, 1)$$

as $n \rightarrow \infty$. A second application of the converging-together principle then yields

$$\sqrt{n} B_n (\delta_i(n) - \partial_i \alpha(\theta_0)) \Rightarrow E[B_1] \sigma N(0, 1).$$

One final application of the converging-together principle (note that $B_n \xrightarrow{a.s.} E[B_1]$ as $n \rightarrow \infty$) proves the theorem. ■

REFERENCES

- Asmussen, S. 1987. *Applied Probability and Queues*, Wiley.
- Bratley, P., B. L. Fox and L. E. Schrage. 1987. *A Guide to Simulation*, Springer-Verlag, New York, second edition.
- Fox, B. L. and P. W. Glynn. 1989. Simulating Discounted Costs. *Management Science*, 35, 1297-1315.
- Glasserman, P. 1991. *Gradient Estimation via Perturbation Analysis*, Kluwer Academic.
- Glynn, P. W. 1990. Likelihood Ratio Gradient Estimation for Stochastic Systems. *Communications of the ACM*, 33, 10, 75-84.
- Glynn, P. W. and Heidelberger, P. 1991. Jackknifing Under a Budget Constraint. *ORSA Journal on Computing*, To appear.
- Goyal, A., P. Shahabuddin, Heidelberger, P., Nicola, V. F., and Glynn, P. W. 1991. A Unified Framework for Simulating Markovian Models of Highly Dependable Systems. *IEEE Transactions on Computers*. To appear.
- Heidelberger, P., X.-R. Cao, M. A. Zazanis, and R. Suri. 1988. Convergence Properties of Infinitesimal Perturbation Analysis Estimates. *Management Science*, 34, 11, 1281-1302.
- Iglehart, D. L. 1975. Simulating Stable Stochastic Systems, V: Comparison of Ratio Estimators. *Naval Research Logistics Quarterly*, 22, 553-565.
- L'Ecuyer, P. 1990. A Unified Version of the IPA, SF, and LR Gradient Estimation Techniques. *Management Sciences*, 36, 11, pp. 1364-1383.
- L'Ecuyer, P. 1991. An Overview of Derivative Estimation. In these proceedings.

- L'Ecuyer, P., N. Giroux, and P. W. Glynn. 1991. Stochastic Optimization by Simulation: Convergence Proofs and Experimental Results for the GI/G/1 Queue in Steady-State. In preparation.
- Miller, R. G. 1974. The Jackknife—A Review. *Biometrika*, 61, 1-15.
- Reiman, M. I. and Weiss, A. 1989. Sensitivity Analysis for Simulations via Likelihood Ratios. *Op. Res.*, 37, 5, pp. 830-844.
- Wolff, R. 1989. *Stochastic Modeling and the Theory of Queues*, Prentice-Hall.

AUTHOR BIOGRAPHIES

PETER W. GLYNN is an associate professor in the Department of Operations Research at Stanford University, Stanford, California. He received the Ph.D. degree from Stanford University in 1982. From 1982 to 1987, he was an assistant professor in the Department of Industrial Engineering at the University of Wisconsin-Madison. His research interests include stochastic systems, computational probability, and simulation. He is an Associate Editor for *Management Science*, *Stochastic Models*, and *Journal of Discrete Event Systems*.

PIERRE L'ECUYER is an associate professor in the department of "Informatique et Recherche Opérationnelle" (IRO), at the University of Montreal. He is also associated with the GERAD research group. He received a Ph.D. in operations research in 1983, from the University of Montreal. From 1983 to 1990, he was a professor in the computer science department, at Laval University, Ste-Foy, Québec. His research interests are in Markov renewal decision processes, sensitivity analysis and optimization of discrete-event stochastic systems, random number generation, and discrete-event simulation in general. He is an Associate Editor for *Management Science* and for the *ACM Transactions on Modeling and Computer Simulation*. He is also a member of ACM, IEEE, ORSA and SCS.

MICHEL ADÈS is associated with the University of Québec at Montréal (UQAM), where he is teaching probability and statistics in the Department of Mathematics and Computer Science. He is also a Ph.D. student in Electrical Engineering at McGill University and is associated with the GERAD research group. His research interests are in stochastic processes (where he published several papers), simulation, and optimal control.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE August 1991		3. REPORT TYPE AND DATES COVERED Technical	
4. TITLE AND SUBTITLE Gradient Estimation for Ratios				5. FUNDING NUMBERS	
6. AUTHOR(S) Peter W. Glynn, Pierre L'Ecuyer, Michel Ades					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of Operations Research - 4022 Terman Engineering Center Stanford University Stanford, CA 94305-4022				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U. S. Army Research Office P. O. Box 12211 Research Triangle Park, NC 27709-2211				10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The view, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation.					
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The ratio estimation problem arises in many different applications settings. This paper is concerned with the interplay between gradient estimation and ratio estimation. Given unbiased estimators for the numerator and the denominator of the ratio, as well as their gradients, joint central-limit theorems for the ratio and its gradient are derived. The resulting confidence regions are of potential interest when optimizing such ratios numerically, or for sensitivity analysis with respect to parameters whose exact value is unknown. The paper also briefly discusses low-bias estimation for the gradient of a ratio.					
14. SUBJECT TERMS ratio estimation, gradient estimation, likelihood ratios, perturbation analysis, simulation				15. NUMBER OF PAGES 8	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL		